

WHAT IS CLAIMED IS:

1. A communication network incorporating a router system, said router system comprising:

an active processor containing a first copy of a router software application that has an active application socket operable to transmit and receive data to and from another network
5 element; and

a backup processor interconnected with said active processor, said backup processor containing a second copy of said router software application that has a replica of said active application socket, said replica application socket operable to transmit data to and receive data from said other network element,

10 such that said replica of said active application socket can maintain an open network connection enabling said backup processor to function as said active processor in the event of failure of said active processor, without the necessity of reestablishing a connection with said other network element.

2. The network of claim 1 wherein said active application socket and said replica application socket are operable to use Transmission Control Protocol (TCP) to communicate with said other network element.

3. The network of claim 1 wherein said active application socket and said replica application socket are operable to use Internet Protocol (IP) to communicate with said other network element.

4. The network of claim 1 wherein said active application socket and said replica application socket are operable to use User Datagram Protocol (UDP) to communicate with said other network element.

5. The network of claim 1 wherein said router software application is a Dynamic Routing Protocol (DRP) application.

6. The network of claim 1 comprising:

at least two said active application sockets each operable to connect with another network element; and

a number of said replica application sockets equal to the number of said active

5 application sockets, each active application socket paired with one said replica application socket, such that each connection established by an active application socket can be maintained in an open and operational state should said backup processor function as said active processor if said active processor fails.

7. The network of claim 1 wherein said router system further comprises a plurality of packet forwarding modules (PFMs), said PFMs being distributed across said router system, each said PFM being interconnected remotely with both of said active processor and said backup processor.

8. The network of claim 7 wherein said PFMs are interconnected with both of said active processor and said backup processor through redundant data links.

9. The network of claim 7 further comprising a plurality of peer routers remote from said router system, said peer routers being interconnected with said router system through said PFMs.

10. The network of claim 1 wherein said backup processor and said active processor are interconnected with one another through a gigabit Ethernet link.

11. A method of connection protection switching in a communication network, comprising:

running concurrently within a router first and second copies of a router software application;

5 creating a first application socket associated with said first copy of said router software application;

creating a replica socket associated with said second copy of said router software application, said replica socket replicating said first application socket;

10 synchronizing state changes between said first and second copies of said router software application;

distributing network route information from said first copy of said router software application to a plurality of packet forwarding modules (PFMs);

15 performing network routing functions between peer routers and said plurality of PFMs using said distributed network route information without direct involvement of said first and second copies of said router software application;

reading incoming messages redundantly at both of said first and second copies of said router software application; and

20 switching over the functions of said first copy to said second copy of said router software application, such that network connections between said peer routers and said plurality of PFMs operate seamlessly without interruption throughout said switching over, without the necessity of reestablishing said network connections.

12. The method of claim 11 wherein said synchronizing comprises transferring explicit state messages between said first and second copies of said router software application.

13. The method of claim 11 wherein said synchronizing comprises maintaining a transaction log at said second copy of said router software application.

14. The method of claim 11 further comprising processing said incoming messages and generating response messages at said first copy of said router software application.

15. The method of claim 14 further comprising processing said incoming messages at said second copy of said router software application and discarding any said response messages generated at said second copy of said router software application.

16. The method of claim 14 wherein said response messages are distributed to said peer routers through said plurality of packet forwarding modules.

17. The method of claim 16 wherein said first and second copies of said router software application communicate with said peer routers using Transmission Control Protocol (TCP).

18. The method of claim 17 further comprising replicating and storing redundant copies of said response messages at both of said first and second copies of said router software application prior to said distributing to said peer routers

5 19. The method of claim 18 wherein:

if an acknowledgement of receipt by a peer router of said distributed response message is received from said peer router within a predetermined timing interval after said distributing, then deleting said stored copies of said distributed response message; and otherwise

10 retransmitting said stored response message.

20. The method of claim 19 wherein said acknowledgement includes a Transmission Control Protocol (TCP) Acknowledgment Number.

21. The method of claim 19 wherein said copies of said stored response message are deleted first from said second copy of said router software application and are then deleted from said first copy of said router software application.

22. The method of claim 19 wherein:

if both first and second copies of said router software application are functioning normally, then retransmitting said stored response message from said first copy of said router software application; and otherwise

5 if said functions of said first copy have been switched over to said second copy of said router software application, then retransmitting said stored response message from said second copy of said router software application.

23. The method of claim 11 wherein said routing software application is a version of Dynamic Routing Protocol (DRP).

24. The method of claim 23 wherein:

if both of said first and second router software application copies are functioning
5 normally, then routing said incoming messages from said PFM first to said replica socket, generating copies of said incoming messages at said replica socket, routing said message copies to said first application socket, generating response messages at said first application socket, and routing said response messages sequentially to said replica socket and out through said PFM; otherwise

10 if said second router software application fails to function and said first router software application copy functions normally, then rerouting said incoming messages from said PFM first to said first application socket, generating response messages at said first application socket, and routing said response messages out through said PFM, such that said messages bypass said replica socket; and otherwise

15 if said first router software application copy fails to function and said second router software application copy functions normally, then rerouting said input messages from said PFM first to said replica socket, generating response messages at said replica socket, and routing said response messages out through said PFM, such that said messages bypass said first application socket and such that the functions of said first router software application
20 copy are switched over to said second router software application copy.

25. The method of claim 24 wherein said rerouting is performed using an action selected from the group consisting of changing Internet Protocol (IP) addresses, aliasing IP addresses, and reprogramming media access controller (MAC) addresses.

26. The method of claim 24 wherein said network connections between said plurality of packet forwarding modules and said peer routers are maintained uninterrupted with at least 90 percent availability throughout said switching over, without the necessity for reestablishing said network connection.

27. The method of claim 24 wherein said network connections between said plurality of packet forwarding modules and said peer routers that are not in process of being established or terminated during said switching over are maintained uninterrupted with at least 90 percent availability throughout said switching over, without the necessity for
5 reestablishing said network connection.

28. The method of claim 11 further comprising splicing new messages following old messages subsequent to said switching over.

29. The method of claim 28 wherein said first router software application copy transmits information that defines message boundaries to said second router software application copy, such that in the event said second router software application copy functions in place of said first router software application copy, said peer routers receive no
5 partial messages.

30. The method of claim 11 wherein said first and second copies of said router software application communicate with said peer routers using Border Gateway Protocol (BGP).

31. The method of claim 11 wherein said first and second copies of said router software application communicate with said peer routers using User Datagram Protocol (UDP) over Internet Protocol (IP).